Collapsing words and worlds; ontological plurality and structural reasoning

Benjamin James

January 31, 2026

We keep learning the hard way that using the same words doesn't necessarily mean we're pointing at the same thing. Two people will say "truth," "harm," "freedom," "safety," "trauma," "intelligence," and feel aligned for a moment, but then their conversation snaps and turns hostile. When you look closely, the turn isn't mysterious. It happens because their words were doing social work, not referential work. They created the feeling of agreement without the substance of shared reference. Their terms functioned like a flag instead of a pointer. This has become everyday life. Agreement by vocabulary, disagreement by reality.

You see it in the simplest conversations. Ask someone what "misinformation" means and you'll get different answers that aren't just differences of opinion; they're differences of category. One person means "factually false statements." Another means "claims that undermine trust." Someone else means "unapproved narratives." And still, another means "speech that leads to harmful outcomes." These are not variations inside one concept. They are different concepts sharing one label. People fight because we believe we are arguing about a claim, when we are actually arguing about what counts as a claim and what counts as proof. Words are shared, but their target is not.

This problem isn't new, but it has changed in character. In the past, most communities lived inside a smaller set of shared assumptions. People still misunderstood each other, but there were fewer world models, the media environment was slower, and the cost of producing convincing narratives was higher. Today we have ontological plurality, where many different ways of carving reality remain viable at once. That plurality is not just cultural diversity or ideological diversity. It is something more basic. It is the fact that people now occupy different "what is real?" templates while speaking a common language. The same sentence can be said with radically different hidden commitments about what exists, what matters, what causes what, and what counts as evidence.

Artificial intelligence amplifies this by making fluent language cheap. A persuasive explanation is no longer a signal that someone has done the work of thinking, and a coherent paragraph is no longer a signal that it was earned by contact with anything checkable. Fluency used to be correlated, however imperfectly, with effort and competence. Now fluency is a commodity, and that changes our ecology of meaning. It doesn't just increase falsehoods; it increases the number of plausible stories a person can live inside without noticing that their story has slipped loose from reality. It accelerates a process that we already struggle with, mistaking a well-formed description for a true one.

The result is a condition where our main threat is not disagreement, but false agreement. People nod along to the same phrases while building incompatible worlds underneath them. They coordinate based on slogans and then discover later that they were never coordinating the same action, using the same evidence, or considering the same constraints. You see this all the time in organizations. Everyone agrees the goal is "accountability," then the work begins and half the room means transparency, the other half means punishment, and the other half means the appearance of process. Or everyone agrees the goal is "equity," then one person means equal outcomes, another means equal opportunity, another means repair for historic injustice, and another means group proportionality enforced by policy. Each meaning implies different tradeoffs. The "agreement" was only lexical, and the bill arrives later.

At this point it is tempting to reach for a grand theory. It is tempting to think we need better epistemology, new metaphysics, or systems thinking, but if we're trying to do something useful, we should start closer to the ground. The most honest thing to admit is that none of us reason outside an ontology. We don't stand on a neutral platform and survey reality from nowhere. We perceive through models, interpret through categories, and compress our experience into language. Even our best attempts at structural reasoning happen inside our own way of carving the world, and that carving changes as we encounter what resists us. The moment you actually engage the real, you pluralize your ontology, because you're forced to revise it. That is learning. The problem is not that we have models. The problem is that we forget we have models, and we start treating our model's vocabulary as the world itself.

So, the goal cannot be perfect objectivity. Our goal is more modest and more urgent, to reduce the distance between our words and what can be directly sensed, jointly observed, or repeatedly tested. The closer we stay to that layer, the more likely we are to be pointing at the same thing. The further we drift into stacked abstractions, the more likely we are to be trading caricatures.

That caricature stack is the dominant failure mode of modern discourse, and you can observe it everywhere once you know what to look for. Something happens in the world. It is specific. It has a time, place, people involved, and messy texture. Then someone describes it with a concept. The concept may be helpful at first; it compresses complexity so we can talk about it. Then the concept gets repeated without the original event. The repetition becomes a new object. Another person repeats the repetition and adds a moral conclusion. Someone else turns the moral conclusion into a slogan. The slogan spreads because it is portable and emotionally legible. Soon, the slogan becomes the thing people argue about. By then, the original event is gone. The discussion has become a loop of symbols referencing symbols. The loop is coherent, and it feels like truth, but it is increasingly untethered from reality.

This is how “issues” are manufactured. Not by conspiracy, but by a predictable compression dynamic. The most portable version of a story survives, and portability is inversely related to accuracy. The portable version loses caveats, context, scale, and the parts that don’t fit the moral narrative. Then the portable version becomes the reference for the next portable version. A caricature becomes the base layer for a new caricature, and then for another. The system produces copies of copies. People call this “public conversation,” but it is closer to a self-feeding symbol machine.

You can tell when a conversation has entered caricature propagation because concrete anchors vanish. Ask for a specific example and the speaker responds with another abstraction. Ask what would change their mind and they treat the question as an insult. Ask what they observed directly and they offer a summary of what “everyone knows.” The content becomes strangely immune to reality, because reality has been replaced by a narrative that now functions as their identity. When someone points out a counterexample, it does not

revise their model; it gets reinterpreted until it fits. Their story becomes unfalsifiable, and unfalsifiable stories are not descriptions. They are social artifacts.

The reason this is so corrosive is that abstractions do not merely describe. They shape what we can imagine. Every abstraction implies a set of interventions, and every model carries a menu of actions and a theory of what will happen next. When our model is a stack of caricatures, our interventions become similarly detached. We start building policies to solve the slogan rather than the underlying situation. We start rewarding compliance with our narrative rather than contact with outcomes. We substitute confidence and fluency for the messy work of adapting to what is actually happening.

This is where "structural reasoning" becomes necessary, and it should be defined in a way that is visible, not theoretical. Structural reasoning is what you do when you refuse to let words float. It is what you do when you keep asking, "What are we actually pointing at?" and "What is this claim resting on?" It is tracing dependencies and constraints rather than trading labels. It is keeping a clean boundary between what was observed and what was inferred. It is maintaining the humility that your model is provisional and having the discipline to update it when reality pushes back.

Notice what this does not require. It does not require you to become a philosopher. It does not require new jargon. It requires a shift in what you treat as acceptable conversation. It requires you to stop rewarding statements for sounding right and start rewarding them for being tethered. It requires you to treat the ability to name a concrete example as a sign of seriousness, and the inability to do so as a sign that the speaker is operating in a closed loop of symbols.

In practical terms, structural reasoning begins with a reference handshake. If someone uses a loaded term, you don't argue the term. You ask what they mean in a way that forces contact with the world. You ask for boundaries and counterexamples. You ask what does not count. You ask what they personally saw, what data they would accept, or what event would make them revise their view. This is not a "gotcha" moment; it is an attempt at repair. It is how you keep the conversation from sliding into mutual hallucination. Most conflicts persist because

this handshake never happens. People argue endlessly inside mismatched ontologies while believing they share a topic.

Structural reasoning also involves naming constraints. Every model has non-negotiables, whether the person admits it or not. One person's model cannot violate autonomy. Another person's model cannot violate stability. Someone else's model cannot violate purity. Yet another individual's model cannot violate compassion. These constraints are not merely preferences; they shape what counts as an acceptable answer. If you don't name them, you end up having moral arguments that are really constraint conflicts. The debate becomes endless because both speakers are optimizing for different invariants while pretending they're solving the same problem. Naming constraints doesn't solve the conflict, but it makes the conflict visible. It turns "you're wrong" into "we are prioritizing different failure modes," and that is a step toward reality.

You can see the absence of structural reasoning in how we treat complex systems. In mental health, people often talk past each other because they use the same diagnostic words while pointing at different phenomena. One person uses "depression" to mean a biochemical condition, while others use it to mean a response to life circumstances, a spiritual crisis, or society's failure to provide meaning. Our arguments become fierce because each frame implies different interventions, and because institutions prefer frames that keep problems simple, legible, and manageable. A label can relieve suffering, but it can also close the inquiry prematurely. Structural reasoning doesn't deny biology. It asks whether the label is tethered to observable patterns and whether the intervention improves the lived conditions that produced the distress. It refuses to treat a word as a solution.

In politics, the absence of structural reasoning is visible in the way conversations become moral theater. People repeat slogans because slogans are legible. They win applause and signal belonging, but slogans do not specify mechanisms. They do not name constraints, they do not track second-order effects, and they do not preserve options. Structural reasoning forces the conversation back to mechanisms and tradeoffs; what incentive changes, what feedback loops reinforce the new behavior, what new failure modes we

create by fixing this one. That is not a luxury. It is the only way to make policy in a world that does not care about our moral clarity.

AI makes this more urgent because it already fills our discourse with “structure-shaped language” that is untethered. It produces talk of systems and feedback loops while remaining indifferent to what can be checked. It can give you a polished causal story for almost anything, which means causal stories become less trustworthy as signals of understanding. Our only defense is to demand commitments that can be tested and to keep bringing the conversation back to observable anchors. What happened, what would count as evidence, what would falsify this, what constraint are we respecting, what tradeoff are we accepting? If you don’t, AI will happily help you build ever more elegant caricatures, and the elegance will feel like truth while reality drifts further away.

None of this means we should, or can, abandon abstractions. Abstractions are how humans think. They are compressions that let us coordinate beyond our immediate experience. However, abstractions are risky technology. They degrade when they are copied without contact with constraint. They become myths that enforce themselves and turn into self-sealing frames. The responsible use of abstraction is to treat it as a provisional compression and to keep it accountable to what it claims to compress. When an abstraction stops tracking what the world does, you downgrade it. You do not defend it as if it were sacred.

This is where the tone of our era needs to change. We need less reverence for fluency and more respect for anchoring. We need to stop treating “confidence” as a proxy for accuracy. We need to stop rewarding people for being able to narrate anything and start rewarding them for being able to say, plainly, “Here’s what I saw, here’s what I inferred, and here’s what would change my mind.” That is not weakness. It is competence under plurality. It is how adults coordinate when they no longer share a single worldview by default.

If this sounds austere, it is because the alternative is already visible. The alternative is a world we already inhabit, where agreement is theatrical and coordination fails routinely. We inhabit a world of sealed ontologies reinforced by infinite narrative supply, where every frame can find its own supporting evidence because evidence has become a rhetorical style rather than a shared constraint. It is a world where institutions respond to instability by enforcing

symbolic conformity, and where dissent is treated as epistemic sin because our system cannot tolerate frame revision. The question is whether we will build a culture that can admit plurality without dissolving into incoherence.

Structural reasoning does not give us a god's-eye truth. It cannot. It happens inside human models that are distorted by human perception and constrained by human limits, but it can do something vital. It can reduce how often we mistake words for pointers. It can keep abstractions accountable and slow the breeding of caricatures. It can preserve option space that collapses when we buy emotional closure at the cost of contact with reality. It can make conversations less about winning and more about finding where our models diverge, which constraints we're prioritizing, and what reality will allow. In a high-plurality, AI-saturated era, structural reasoning is survival equipment.